

Uncertainty, Redistribution, and the Labor Market^{*}

by Casey B. Mulligan

University of Chicago

February 2013

Abstract

Uncertainty and its composition can affect the demand for social insurance, and thereby the labor market. This paper shows that small to medium-sized increases in uncertainty or risk aversion are enough to recommend an expansion of the safety net that would be broadly similar to the actual safety net expansions, which significantly depressed the labor market. Labor market effects of uncertainty through investment and insurance channels are also examined with employer and employee labor wedges.

^{*} I appreciate comments from Gianluca Violante and seminar participants at the University of Chicago, and the financial support of the George J. Stigler Center for the Study of the Economy and the State.

During the depths of the Great Depression, Franklin Delano Roosevelt famously said “The only thing we have to fear is fear itself.” (White House 2012) Americans and even economists were fearful during the fall of 2008 when large financial firms went bankrupt and financial markets were unable to operate normally. The public sector has been accused of creating significant uncertainty during the past couple of years by ceasing normal budgeting procedures, creating an extraordinary volume of new regulation for the health and financial sectors, and running up the national debt with no clear plan as to paying it back.

Various measures of uncertainty have emerged (Baker, Bloom and Davis 2011), as have explanations of why labor market activity might be sensitive to the amount of uncertainty (Bloom, et al. 2011). Even without drawing a distinction between fear and uncertainty, the explanations are incomplete because they tend to be investment intensive, and do not address the fact that wages and the quantity of labor moved in opposite directions over the past couple of business cycles.¹ On the first point, a familiar story is that uncertainty makes firms unwilling to invest, or their creditors willing to lend the funds to invest, and that the makers of investment goods find themselves without work. Yet real investment per capita in, say, 2011 was \$2,000 (2005 prices) less than it was in 2007, before the recession began. \$2,000 per capita is only

¹ More specifically, models such as Bloom, et al. (2011, see especially their equation 12) fail to generate the empirically observed “labor wedge” – a recession in which consumption and labor both fall while real wages do not. The body of this paper further discusses the labor wedge. More work is also needed to examine whether uncertainty *increases* investment, and labor more directly, through precautionary savings motives (Carroll 1997) .

about 4 percent of GDP, and a 4 percent reduction in aggregate labor demand is expected to reduce the aggregate quantity of labor by about one percent.² But in fact the employment-population ratio was a full 7 percent less in 2011 than it was in 2007. In other words, even if we assume that the entire investment drop was a consequence of fear or uncertainty, the vast majority of the decline in the quantity of labor remains unexplained unless fear and uncertainty affect the labor market through another mechanism.

At the same time that fear and uncertainty spiked and the labor market contracted, various subsidy programs for the poor and unemployed expanded their eligibility rules and implemented more generous benefit formulas. Using the principal-agent model of the “equity-efficiency tradeoff”, this paper suggests that the demand for social insurance is an important, if not the primary, link between uncertainty and the quantity of labor. Before the recession, uncertainty and fear were relatively low, and society had settled on a balance between the amount of redistribution or social insurance and the amount of labor market inefficiency in the forms of a distorted size and composition of the workforce. This pre-recession balance was a function of both the technological tradeoff between equity and efficiency, the degree of risk aversion, and political factors. With the greater amounts of uncertainty and fear that prevailed after 2007, the technological tradeoff was different and/or risk aversion was greater, and society changed the balance between equity and efficiency in the rational direction of less labor market efficiency.

The size of this effect depends on the magnitude of the changes in social insurance program rules, which I have quantified elsewhere in terms of implicit marginal labor income tax

² Here I assume a wage elasticity of aggregate labor demand of -3.3 (as it would be with an aggregate Cobb-Douglas production function with labor share 0.7) and a wage elasticity of aggregate labor supply of one. The labor quantity impact would be even less if the wage elasticity of labor supply were less than one, or investment industries were more capital-intensive than the rest of the economy.

rate changes (Mulligan 2012). Statutory marginal labor income tax rates facing the typical household head or spouse at the middle of the skill distribution increased about eight percentage points, and thereby reduced the after-tax return to working by almost fifteen percent, in less than two years. Marginal tax rate changes of this magnitude can easily reduce the quantity of labor by several percentage points.

A Model of the Equity-Efficiency Tradeoff

In the principal-agent model of the labor market, workers devote time and effort n to production. This includes time spent at work, time spent searching for new work, and effort devoted to enhancing and maintaining the productivity of those uses of time. The value produced by that effort is $y = n + v + \varepsilon$, where v and ε are the result of mean zero idiosyncratic random factors beyond the worker's control. For simplicity, the random factors act additively on the worker's effort, rather than multiplicatively as in Mirrlees (1971). My additive model is just a small adaptation of a special case of Holmstrom and Milgrom (1987) that has been applied to executive salaries by Rosen (1982), Garen (1994), and others.

Some of the random factors, embodied in v , are widely observable. The others, embodied in ε , are not observed, except by the worker himself who can infer their combined value by subtracting his effort and v from the value produced. I normalize ε so that it is uncorrelated with v , and assume that the worker is risk-averse (more on this below). As a result, the worker optimally pools the result v of observed idiosyncratic random factors in a full-information insurance market: when v is observed he pays it to the insurance group, or receives $-v$ from the

group in case its value is negative. In theory the insurance group is infinitely large and free from administrative costs so that total insurance premiums (from the group members with positive v s) exactly finance total insurance awards (paid to group members with negative v s). In practice, the insurance group may be co-workers, family members, church members, etc.

Even with insurance against the observed random factors, the remaining value produced $n + \varepsilon$ is still random. The purpose of this paper is to examine possible ways of insuring the idiosyncratic ε risk, and deriving the effects of the amount and composition of uncertainty on the efficient amount of that insurance.

Suppose that a worker additionally enters into an imperfect information social insurance arrangement in which she pays a fixed fraction of her earnings $n + \varepsilon$ (net of payments with respect to full-information v insurance) and receives an insurance benefit that is a linear function of net income. Her disposable income c is therefore a linear function of earnings, the intercept of which I denote as b and the slope I denote as μ :

$$c = (n + \varepsilon)\mu + b \quad (1)$$

The social insurance system has a budget constraint that relates benefits it pays to people having zero income to payments it receives from people having positive income:

$$b = \frac{1 - \mu}{1 + \phi} \bar{n} \quad (2)$$

where \bar{n} is the average effort by members of the insurance system. $\phi \geq 0$ is an administrative or stigma cost reflecting the possibility that social insurance benefits might be worth less than they cost. The average idiosyncratic risk ε is zero but I explain below how the cross-section distribution might be related to aggregate “shocks.”

Workers have a smooth von Neumann and Morgenstern (1944) utility function u defined over their value of consumption net of the costs of effort, with strictly positive first derivative and strictly negative second derivative in the relevant range. Their expected utility function is therefore:

$$\int u \left(\mu n + \mu \varepsilon + \frac{1-\mu}{1+\phi} \bar{n} - \gamma \frac{\eta}{\eta+1} n^{(\eta+1)/\eta} \right) dG(\varepsilon) \quad (3)$$

where G is the distribution of random factors ε (continuous, with all finite moments), and the constant $\eta > 0$ is the wage elasticity of labor supply. The constant $\gamma > 0$ is a preference parameter, so that the marginal rate of substitution between labor n and disposable income c has no wealth effects on labor. As a result, any worker facing this social insurance program will desire labor in the amount $n = (\mu/\gamma)^\eta$. The desired labor increases with μ because μ is the reward the worker receives for exerting effort; the remaining fraction $(1-\mu)$ of results of effort goes to the social insurance system.

I assume that actual labor coincides with desired labor, which means that the social insurance system has no influence on the worker's labor decision aside from the parameters μ and b of the safety net benefit formula. As a result, average labor \bar{n} is also equal to $(\mu/\gamma)^\eta$, which is generally different from the socially efficient labor $n^* = (1/\gamma)^\eta$ that equates the marginal rate of substitution to the social product of effort.

Because workers choose their effort before knowing the final outcome ε , the only source of randomness in (3) is the $\mu\varepsilon$ term. A smaller value for μ therefore means a lesser amount of risk faced by the worker, and more equal disposable income for workers who end up with different values of ε , but it also means less effort and thereby less aggregate income. This is known as the “equity-efficiency tradeoff,” or less often as a “safety-efficiency tradeoff.”

The equity-efficiency tradeoff can be analyzed quantitatively by considering a measure of equity or safety S that is negatively related to the standard deviation of disposable income $S = (1 + s_c)^{-1} = (1 + \mu s_\varepsilon)^{-1}$, where s_ε is the standard deviation of the random factor ε . By eliminating μ from the safety equation and the average labor equation $\bar{n} = \mu^\eta n^*$, we have a single equation for the tradeoff:

$$S = \frac{1}{1 + s_\varepsilon \left(\frac{\bar{n}}{n^*} \right)^{1/\eta}} \quad (4)$$

More safety S achieved with more social insurance $(1-\mu)$ means less labor efficiency \bar{n}/n^* , and vice versa.³ Figure 1 displays the tradeoff in a simple diagram, with safety measured on the vertical axis and efficiency on the horizontal axis. Points to the northwest on the frontier shown in Figure 1 correspond to more social insurance—that is, low values for μ —and points to the southeast correspond to less social insurance. The red line is the equity-efficiency frontier, because combinations of equity and labor efficiency beyond it are not possible since equalizing outcomes reduces incentives to supply labor.

Possible Changes in the Equity-Efficiency Tradeoff, and the Optimal Degree of Social Insurance

³ Both safety and efficiency measures vary between zero and 1. Safety is zero and efficiency is 1 when the self-reliance rate is 1 (no sharing of the imperfect information risks). Safety is 1 and efficiency is zero when all of the risks are shared.

Both safety and labor efficiency are desirable, and overall efficiency strikes a balance between the two. The optimal amount of social insurance can be described as the value of μ —a point on the equity-efficiency tradeoff shown in Figure 1—that maximizes worker expected utility.

$$\int u \left(\mu \varepsilon + \frac{1 + \phi \mu}{1 + \phi} (\mu / \gamma)^\eta - \gamma \frac{\eta}{\eta + 1} (\mu / \gamma)^{\eta + 1} \right) dG(\varepsilon) \quad (5)$$

It is straightforward to prove that the optimal amount of social insurance depends on the amount of risk the worker faces, as embodied in the distribution function G .

Proposition 1 The optimal amount of social insurance ($1 - \mu$) increases with the standard deviation s_ε of the random factor ε , holding constant higher-order moments of the distribution G .

Proof Because u is smooth, I prove the proposition by using its Taylor expansion (in the neighborhood of $\varepsilon = 0$) in the worker expected utility expression:

$$\begin{aligned} & u \left(\frac{1 + \phi \mu}{1 + \phi} (\mu / \gamma)^\eta - \gamma \frac{\eta}{\eta + 1} (\mu / \gamma)^{\eta + 1} \right) + \int \sum_{k=2}^{\infty} \frac{u^{(k)}}{k!} (\mu \varepsilon)^k dG(\varepsilon) \\ & = u \left(\frac{1 + \phi \mu}{1 + \phi} (\mu / \gamma)^\eta - \gamma \frac{\eta}{\eta + 1} (\mu / \gamma)^{\eta + 1} \right) + \sum_{k=2}^{\infty} \frac{u^{(k)}}{k!} \mu^k G_k \end{aligned}$$

where $u^{(k)}$ denotes the k th derivative of u evaluated at $\varepsilon = 0$ and G_k denotes the k th moment of G .

The first derivative of expected utility with respect to μ has one term involving the second moment of G , and it is linear with a negative coefficient because u is concave. Thus the optimal μ falls, and the optimal $(1 - \mu)$ increases, with the second moment.

There are a couple of reasons to think that standard deviation s_ε of the random factor ε was greater after 2008 than it was before. One is that capital market and other aggregate events made it more difficult for market participants to distinguish bad outcomes that should be blamed on low effort from bad outcomes that were just unlucky. In this view, an aggregate shock does

not necessarily change the total amount of idiosyncratic risk – the variation of labor productivity y around labor time and effort supplied – but rather changes its composition between the unobserved category ε than the observed category v . When the change is in the direction of more unobserved idiosyncratic risk, workers replace part of the full-information insurance lost by adding to their imperfect-information insurance, despite the latter’s cost in terms of labor inefficiency. Figure 2 illustrates the change that might have occurred between 2007 and 2009. As workers began to face more risk—the equity-efficiency frontier shifted down from red to gray—the safety net could have remained as it was in 2007, in which case labor market efficiency might have remained constant too. But then workers would have substantially less safety, as with the hollow circle in Figure 2. In that unfortunate situation, workers may prefer to recover some of the lost safety by reducing labor efficiency, as with the solid circle on the gray 2009 equity-efficiency frontier.⁴

A related argument is that the variance of ε increased without any reduction in the variance of v . Still, workers desire to have more imperfect-information insurance, despite its cost in terms of labor inefficiency. Either way, less labor is the result of the safety net expansion in the sense that the safety net could have been kept constant, in which case labor would have remained constant. But the safety net expansions were themselves a response to another shock—increased uncertainty—and keeping the safety net constant might not have been the best response to that additional uncertainty. The recession, which is by definition a drop in aggregate labor market activity, may have been preventable by keeping the safety net constant, but the cost of preventing it would have been too high in terms of the risks that workers would have borne

⁴ A related argument is that the equity-efficiency tradeoff is fixed, but the optimal balance of equity and efficiency has changed since 2007 because people became more risk-averse, at least temporarily. See Holmstrom and Milgrom (1987, 323) for a risk-aversion comparative static.

without government help. In this sense, one might say the ultimate cause of the recession was the change in the equity-efficiency tradeoff mediated through an expanded safety net.

A smaller value for the stigma parameter ϕ induces an optimal movement along Figure 1's equity-efficiency frontier in the direction of more social insurance. A greater degree of risk aversion has a similar effect.

Quantitative Estimates of Risk and Marginal Tax Rate Changes since 2007

Economists have noted that uncertainty has increased a lot since 2007, and they have offered measures of the amount of added uncertainty (Bloom, et al. 2011).⁵ One commonly cited measure is the Chicago Board Options Exchange's Volatility Index, sometimes called the fear index or VIX, which is based on market forecasts of the volatility of an index of stock prices. The monthly version of the series is plotted in Figure 3, and has percentage units of thirty-day annualized rates of return (Chicago Board Options Exchange 2009). The index increased by a factor of two from the first half of 2007 to the second half, and then by at least another factor of two by the end of 2008. The index returned to its late-2007 level in early 2010 and again in late 2010.⁶

⁵ So far, economists have not suggested that the labor market effects of uncertainty would work through marginal tax rates, but rather that uncertainty discourages firms from undertaking new investment projects or otherwise expanding their payrolls (Bloom, et al. 2011).

⁶ Note that the VIX is probably not a good indicator of idiosyncratic risk (as opposed to aggregate risk, which depends on whether the volatility of average stock prices derives from redistribution or aggregate shocks), which is featured in my model of the equity-efficiency tradeoff. However, to the extent that

Even if the amount and composition of uncertainty remained constant, the degree of risk aversion could vary over time, which could simultaneously change asset prices (Campbell and Cochrane 1999) and the demand for social insurance (represented in my model as the degree of concavity of the utility function u). For example, the economic expansion prior to 2007 might have been a time of low risk aversion in which people were especially willing to purchase risky assets and especially eager to forgo social insurance in order to have a more efficient labor market.

Monetary economists also say that the demand for liquidity increased following 2007 (Diamond and Rajan 2011), and surging prices of U.S. government securities are consistent with their interpretation (Lucas 2008). Safety net payments may share some liquidity and risk characteristics with government securities—indeed, food stamp program beneficiaries receive their benefits on a debit card (Eslami, Filion and Strayer 2011)—which could be modeled as an increase in the preference for, or a decrease in the preference against, safety net payments. In my model, a reduction in the welfare stigma cost ϕ is another way that the optimal amount of social insurance could increase, and labor efficiency decrease.

Holmstrom and Milgrom (1987, 323) give an example, consistent with the model above, with constant absolute risk aversion in which the optimal self-reliance rate (μ in my notation, α in theirs) has a closed form solution as a function of the standard deviation of the income risk being insured (s_ε in my notation), the coefficient of absolute risk aversion r , and a parameter governing the disutility of effort. Holding the effort-disutility parameter constant, the comparative statics for their model's optimal self-reliance rate are (my notation):

wage or earnings inequality is an indicator of idiosyncratic risk, it suggests that idiosyncratic risk may have increased significantly since 2007.

$$d \ln \mu = -(1 - \mu)2(d \ln s_\varepsilon) - (1 - \mu)(d \ln r) \quad (6)$$

Appendix I contains the derivation of equation (6).

Figure 4 displays Mulligan's (2012) measure of the statutory marginal labor income tax rate for household heads and spouses in the middle of the skill distribution, including the marginal tax rates implicit in the many federal and state anti-poverty programs. Its changes derive primarily from changes in eligibility and benefit rules for the unemployment insurance program and related subsidies, and for the food stamp program. Its level (but not its changes) reflects actual rates of program participation, which is why the rates shown in Figure 4 are less than reported by studies that examine the rates faced by persons who participate in all programs for which they are eligible (Romich, Simmelink and Holt 2007).

The marginal labor income tax rate began at about 40 percent and peaked in late 2009 at more than 48 percent. In terms of the log self-reliance rate, that's a decrease of almost 0.15.

Assuming for the moment that the optimal self-reliance rate were also about 0.44, then the comparative static elasticities in equation (6) are about -0.88 and -0.44, respectively. Thus, an increase in the log of that standard deviation of about 0.17 would justify a reduction in the log self-reliance rate of 0.15. The same self-reliance rate reduction could also be justified by a 0.34 increase in the log of the coefficient of absolute risk aversion, or some combination of added variance and added risk aversion. Moreover, microeconomic wage data suggest that wage variability was high during the recession, with a log standard deviation change of about 0.07 (see Appendix II). Thus, the Holmstrom and Milgrom (1987) approach does not require especially

large increases in uncertainty or risk aversion in order to recommend an expansion of the safety net that would be broadly similar to the actual safety net expansions.⁷

Regardless of whether it was optimal to reduce the log self-reliance rate by 0.15, it happened and that reduction is expected to significantly depress the labor market. Mulligan (2012) finds that marginal tax rate changes of the amounts shown in Figure 4, with a 0.4 wage elasticity of *aggregate* labor supply that is less than most reported in the literature, explain at least half of the actual reduction in per capita work hours between the end of 2007 and the end of 2009. With a larger (and thereby more reasonable) wage elasticity of aggregate labor supply, even more of the hours decline can be explained.

Wages and Wedges in the Labor Market

Labor wedges are another way to categorize various effects of uncertainty and other factors on labor market outcomes, and thereby distinguish some of the channels by which uncertainty might be depressing the labor market. When employers perceive a new cost of employment in addition to employee compensation, or decide to increase their markup over marginal cost, labor productivity grows more than compensation per hour. Equivalently, compensation per hour grows less than labor productivity, because employers hire up to the point that their marginal revenue product of labor equals the total cost of hiring additional labor. A per-employee fine, such as the fines levied by the Affordable Care Act, are examples of new employment costs in addition to employee compensation. In theory, these taxes create a

⁷ This approach is arguably too powerful because it predicts significant increases in the amount of redistribution when they do not occur. See, for example, the Meltzer and Richard (1981) model of inequality and redistribution and Peltzman (1980) and Perotti (1996) on its incongruence with the time series and cross-country evidence.

productivity-wage wedge of the same percentage as the tax itself. Anticipated severance costs, and other costs of downsizing, create the same sort of wedge between contemporaneous wages and productivity, and this wedge would grow as uncertainty increased the probability of downsizing events.

When employees perceive a new cost of employment, such as a new payroll tax on employees, added child-care costs, or new subsidies for the unemployed, the households' marginal rate of substitution (*MRS*, sometimes described as the reservation wage of the marginal worker) falls more, or rises less, than wages do.⁸ As explained above, added uncertainty may motivate more safety net program benefits, which add to this wedge with a percentage equal to the percentage reduction in the post-fisc share of wages accruing to employees as a consequence of their work decision. Together the growth in the employer and employee wedges cause labor productivity to grow more, and fall less, than the *MRS*.

$$\begin{aligned} d \ln(\text{labor productivity}) &= d \ln(\text{employer wedge}) + d \ln(\text{real wage}) \\ d \ln(\text{real wage}) &= d \ln(\text{employee wedge}) + d \ln MRS \end{aligned} \tag{7}$$

Note that both wedges have log changes that are in the same units, which are essentially units of a tax rate. For example, to a first order approximation a two percent payroll tax would increase the log employer wedge by 0.02 if levied on employer (and not counted in employee compensation for the purposes of measuring the real wage) and increase the log employee wedge by 0.02 if levied on employee.

Figure 5 displays the quarterly employer (dashed) and employee (solid) wedges from 2007 through the end of 2011, applying the formula (7) measuring productivity as real GDP per

⁸ The employee wedge captures any wealth effect or precautionary savings motive as long as it is adequately summarized by consumption as it in the permanent income model.

manhour, the real wage as real employee compensation per manhour, and the log *MRS* change as the log change in real consumption per capita plus the log change in hours worked per adult.⁹ Also shown as a dotted series in the Figure is minus one times the log change in the after-tax share formed from Figure 4's marginal tax rate series. In theory, the solid and dotted series should coincide, at least if the marginal tax rate and other series were not measured with error and the proper marginal rate of substitution function were used. Although the two series are constructed from completely different data sources, they coincide closely through 2010 both in terms of direction and magnitude. Both the after-tax share and the employee wedge fall during 2011, but the wedge falls significantly less.

The employer wedge is relatively constant through mid-2009, after which time its log increases almost 0.05 from mid-2009 to the end of 2011 when it is 0.035 greater than it was when the recession began. The employer wedge series is consistent with the hypothesis that *something* reduced the demand for labor (at a given level of real compensation per hour) after mid-2009. That “something” could have been the effects of uncertainty working through labor demand channels.¹⁰ In any case, the changes in the employer wedge are large by historical standards – roughly of the magnitude of a three percentage point increase in the employer payroll tax – but small in comparison to changes in the employee wedge or changes in the after-tax share (they look more like a change in the employee payroll tax exceeding ten percentage points),

⁹ This marginal rate of substitution function corresponds to a representative agent's utility function of the form

$$u(c, n, P, A) = P \ln(c / P) - \gamma \frac{\eta}{\eta + 1} A \left(\frac{n}{A} \right)^{\eta/(\eta+1)}, \text{ where } c \text{ is aggregate real consumption, } n \text{ is aggregate hours}$$

worked, P is the total population, and A is the adult population. The constant η is usually interpreted as the Frisch elasticity of aggregate labor supply with respect to after-tax real wages, and is taken to be one for the purposes of constructing Figure 5's employee wedge. The employee wedge is similar with η equal to say, 0.75, but would increase significantly more if η were taken to be less than 0.5.

¹⁰ A growing employer wedge could also reflect increases in the amount of social insurance via employer-related instruments, such as the employer fines used in the new Affordable Care Act to finance part of its health insurance subsidies.

which is one of the reasons why it is worth considering the idea that uncertainty affects the labor market through labor supply at least as much as it works through labor demand.

Conclusions

Safety net programs face a well-known equity-efficiency tradeoff: providing more resources for the poor can raise their living standards, but it also gives them less incentive to raise their own living standards. Most societies somehow balance the tradeoff by permitting outcomes to vary across people—less-than-perfect equity—yet still having some amount of redistribution.

The recent housing crash, financial crisis, or even political events may have altered the nature of this tradeoff or altered society's willingness to tolerate labor market inefficiency in order to have more equity. The tradeoff would be altered by a deterioration of information about labor market outcomes that would normally help distinguish luck from the consequences of worker effort. It could also be altered by an increase in non-diversifiable risk. An increase in risk aversion, or a reduction in welfare stigma, induces an optimal movement along that tradeoff in the direction of more social insurance. Under any of these scenarios, it would be no surprise that the social safety net expanded and labor market activity fell as much as they did, and uncertainty in one way or another would be the ultimate cause.

A previous literature has emphasized effects of uncertainty on the labor market through labor demand, in the sense that an employer's desired hiring depends not only on what employees will be paid in the short term but also on the degree of uncertainty about their future pay or other future costs of doing business. To the degree that wages adjust over time to reflect supply and demand, more uncertainty should in this view cause real wages to grow less than

productivity, which did not happen in the two or three years prior to mid-2009, but did happen after that. The growth of the employer wedge after mid-2009 amounted to about a 3-5 percent increase in the marginal payroll tax rate on employers, and is large by historical standards.

However, much of what happened in the labor market happened on the employee side. The marginal tax rate on employees, most of it implicit in safety net program spending rules, increased three or four times more than the employer wedge did. To the degree that fear and uncertainty about the economy motivated the new safety net program rules, they affected the labor market more through labor supply than they did through labor demand.

This paper examines the demand for social insurance in the context of an optimal social insurance model. I do not interpret the model results as proving that the actual safety net expansions of recent years were optimal. Indeed, the actual expansions had significant flaws, such as subjecting parts of the population to 100 percent tax rates (Mulligan 2013), subjecting households to additional policy randomness (Baker, Bloom and Davis 2011), and allowing households to participate in multiple uncoordinated insurance schemes (Mulligan 2012). Nevertheless, the model shows why we should not expect the social safety net to remain constant, and thereby not expect the labor market to remain constant, while households deal with added uncertainty.

Appendix I: ARA Closed-Form Solution

Following Holmstrom and Milgrom (1987, 323), assume a normal distribution for ε and let the utility function be of the ARA form $u(x) = -e^{-rx}$, let the wage elasticity of effort supply be one, and let the stigma cost ϕ be zero. Equation (5) becomes

$$\int u \left(\mu \varepsilon + \frac{1+\phi\mu}{1+\phi} (\mu/\gamma)^\eta - \gamma \frac{\eta}{\eta+1} (\mu/\gamma)^{\eta+1} \right) dG(\varepsilon) = -e^{-r \frac{1+\phi\mu}{1+\phi} (\mu/\gamma) + r\gamma (\mu/\gamma)^2/2 + \mu^2 r^2 s_\varepsilon^2/2} \quad (8)$$

where s_ε is the standard deviation of the idiosyncratic private information risk. Maximizing with respect to the self-reliance rate μ ,

$$\mu = (1 + r s_\varepsilon^2 \gamma)^{-1} \quad (9)$$

Holding constant the effort disutility parameter γ , the comparative statics are the same as shown in the main text's equation (6):

$$d \ln \mu = -(1-\mu)2(d \ln s_\varepsilon) - (1-\mu)(d \ln r) \quad (6)$$

Appendix II: Idiosyncratic Risk as Residual Wage Inequality

Two key parameters in the model are changes over time in the risk aversion coefficient and in the amount of idiosyncratic private information risk. Both are difficult to measure directly. Outcome variability can be measured, but it's difficult to know how much outcome variability is the consequence of effort, or the consequence of anticipated factors, or the consequence of shocks whose origins are publicly known and thereby insured without moral hazard. Advancing risk measurement is beyond the scope of this paper, but this appendix offers a crude measure that might indicate the general direction of idiosyncratic private information risk changes and perhaps roughly indicate their magnitude.

One measure is a residual from an hourly low wage regression on indicator variables for white; calendar year; month of year; the interaction of an age quartic with educational attainment and sex; the interaction of educational attainment, sex, and presence of children under eighteen; and all interacted with married (spouse present). The sample was full-time employed persons aged 26-64 in the monthly CPS Merged Outgoing Rotation Groups for 2000 through 2010. The standard deviation of these residuals averaged 0.512 for 2006-2010, as compared to 0.551 for 2009-10: a log change of 0.073.

Figure 1. The Equity-Efficiency Frontier

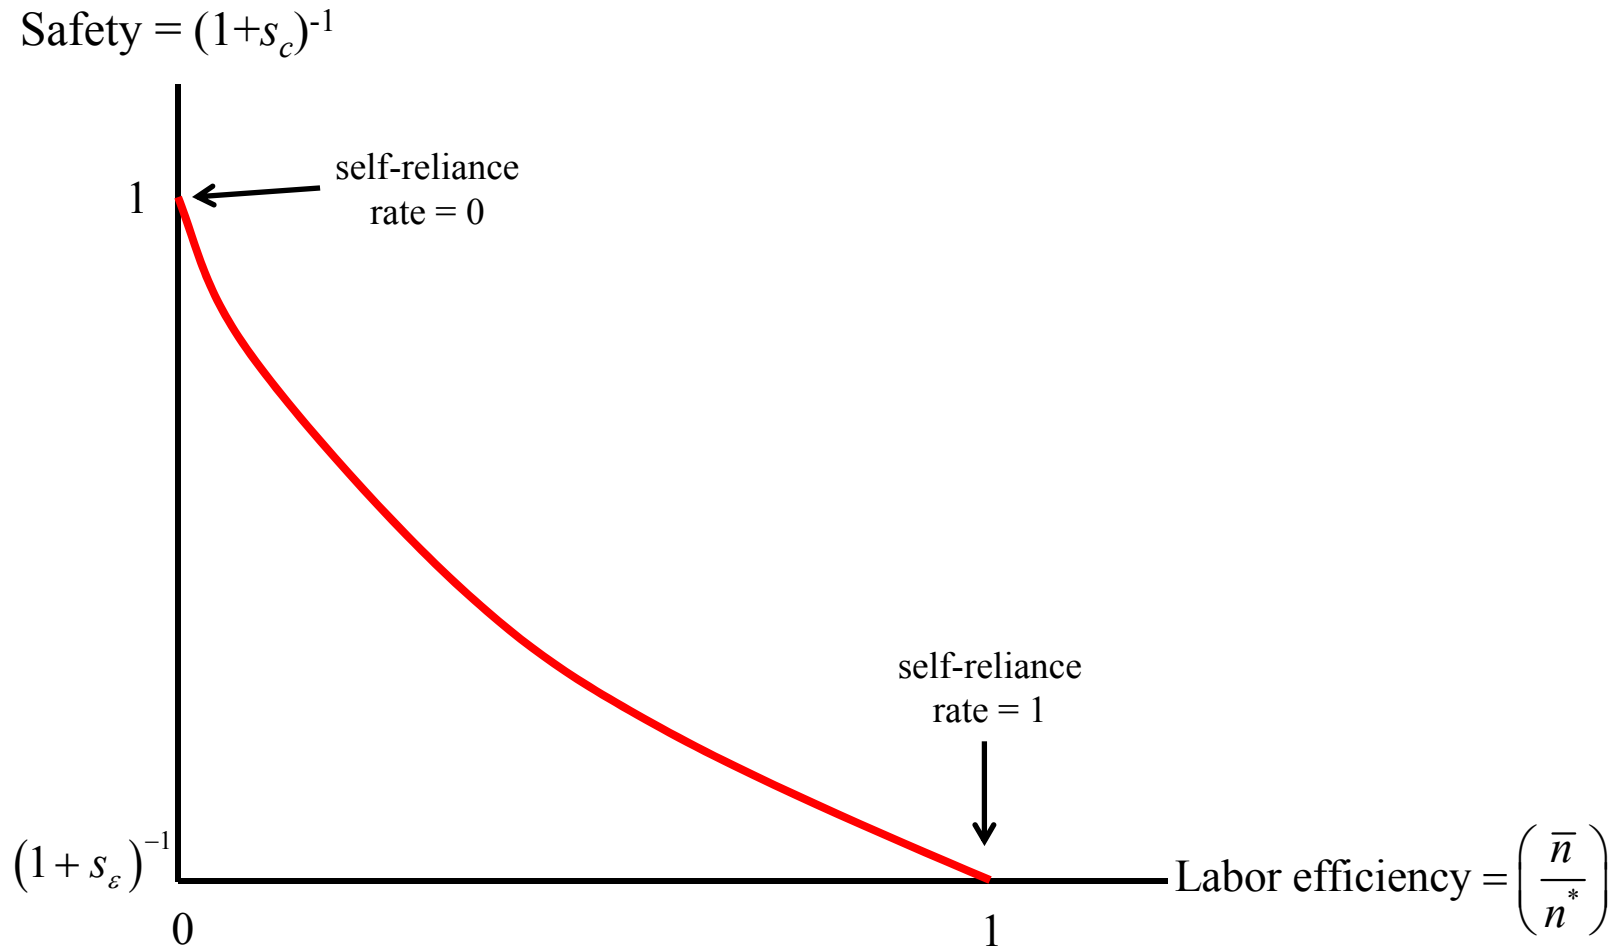


Figure 2. Changes in the Equity-Efficiency Frontier

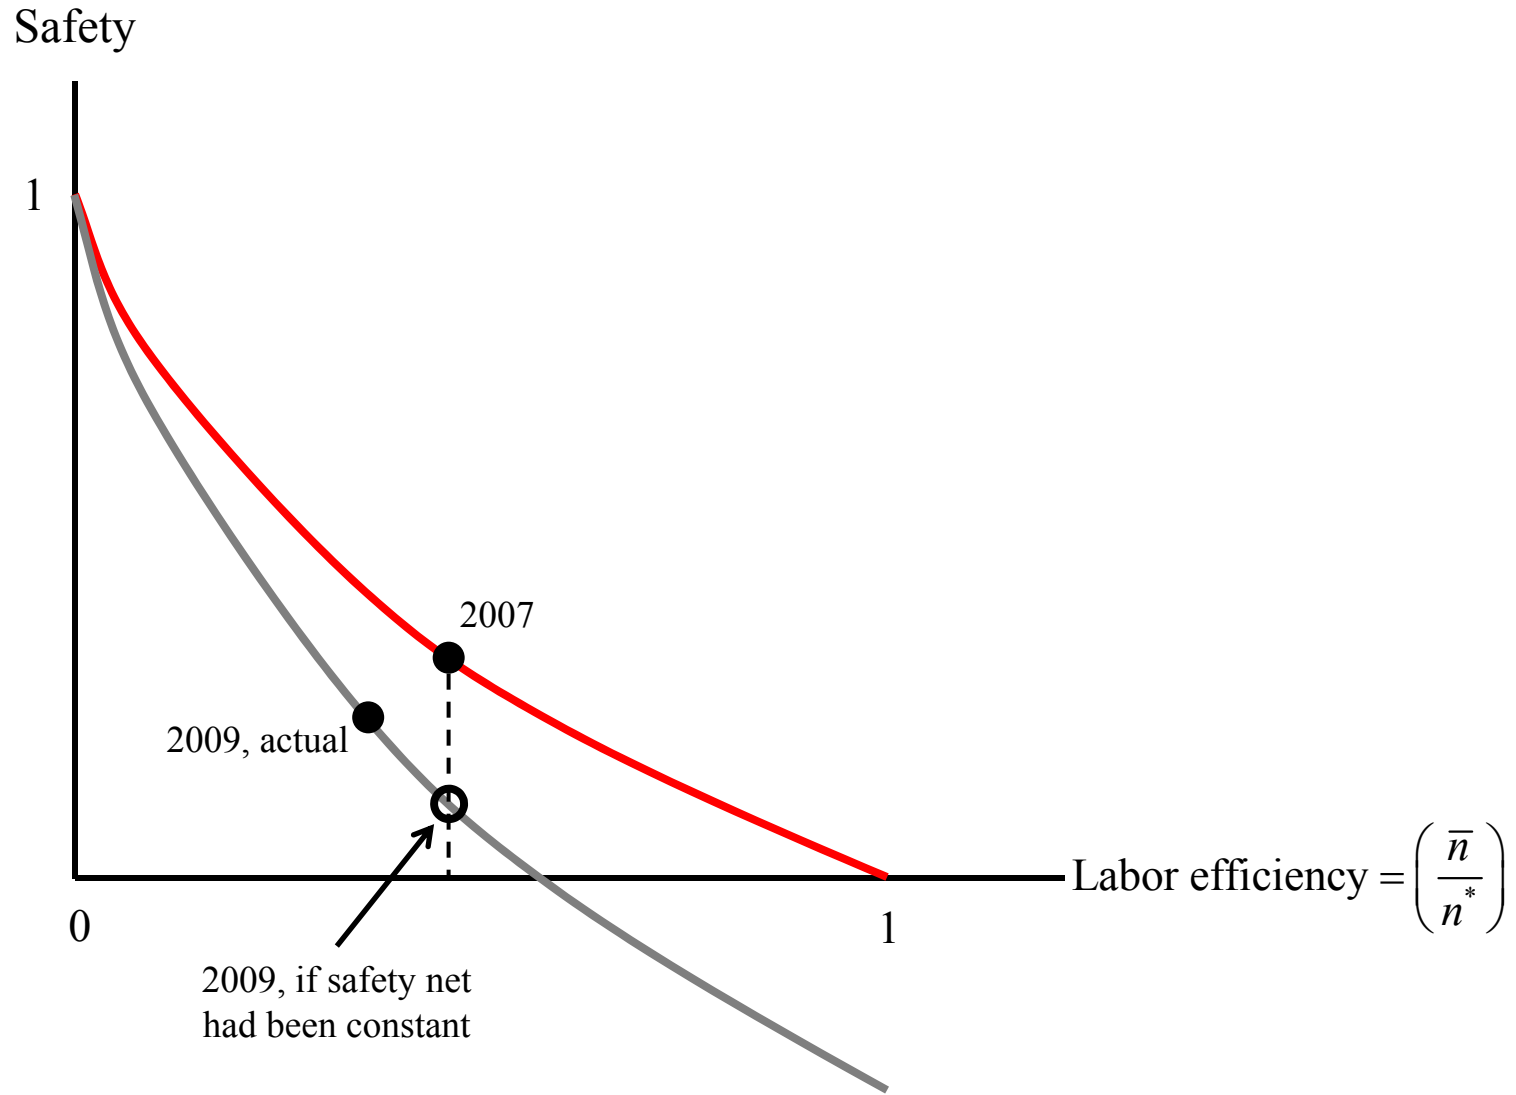


Figure 3. The CBOE Volatility Index, monthly 2006-2010

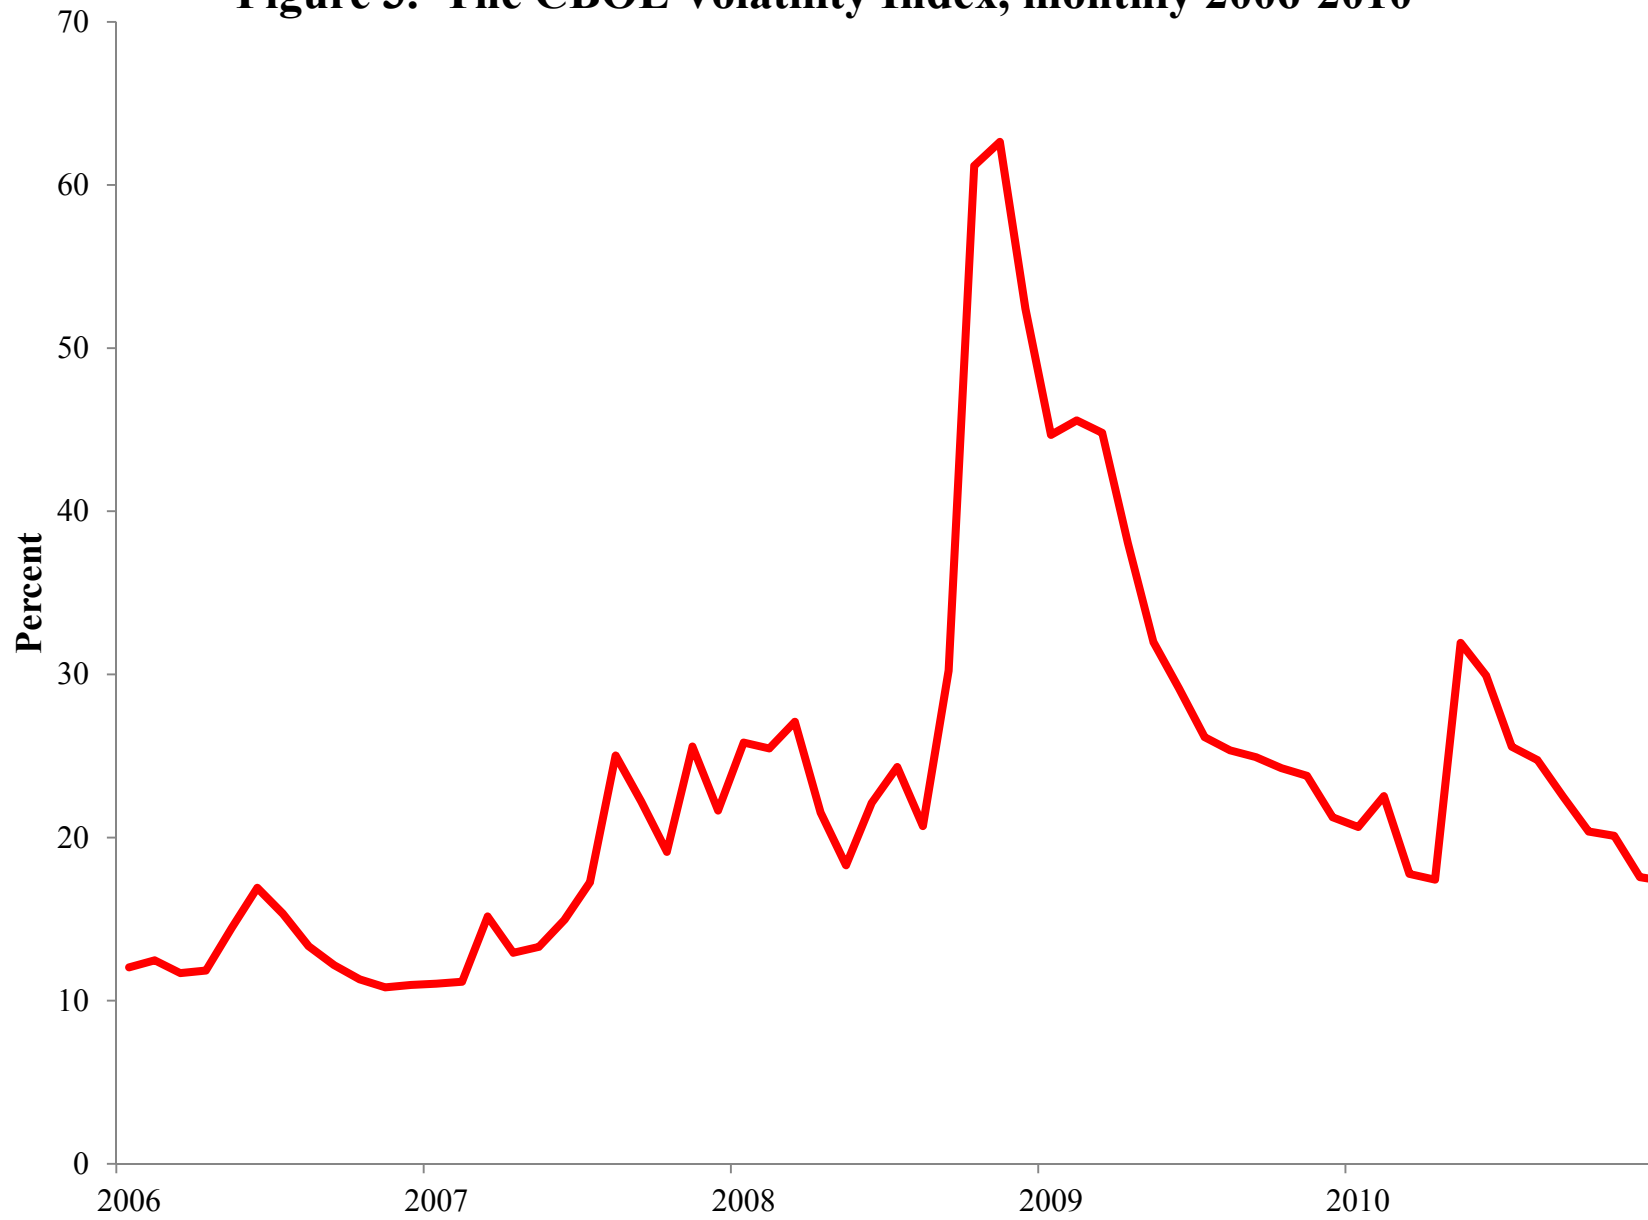
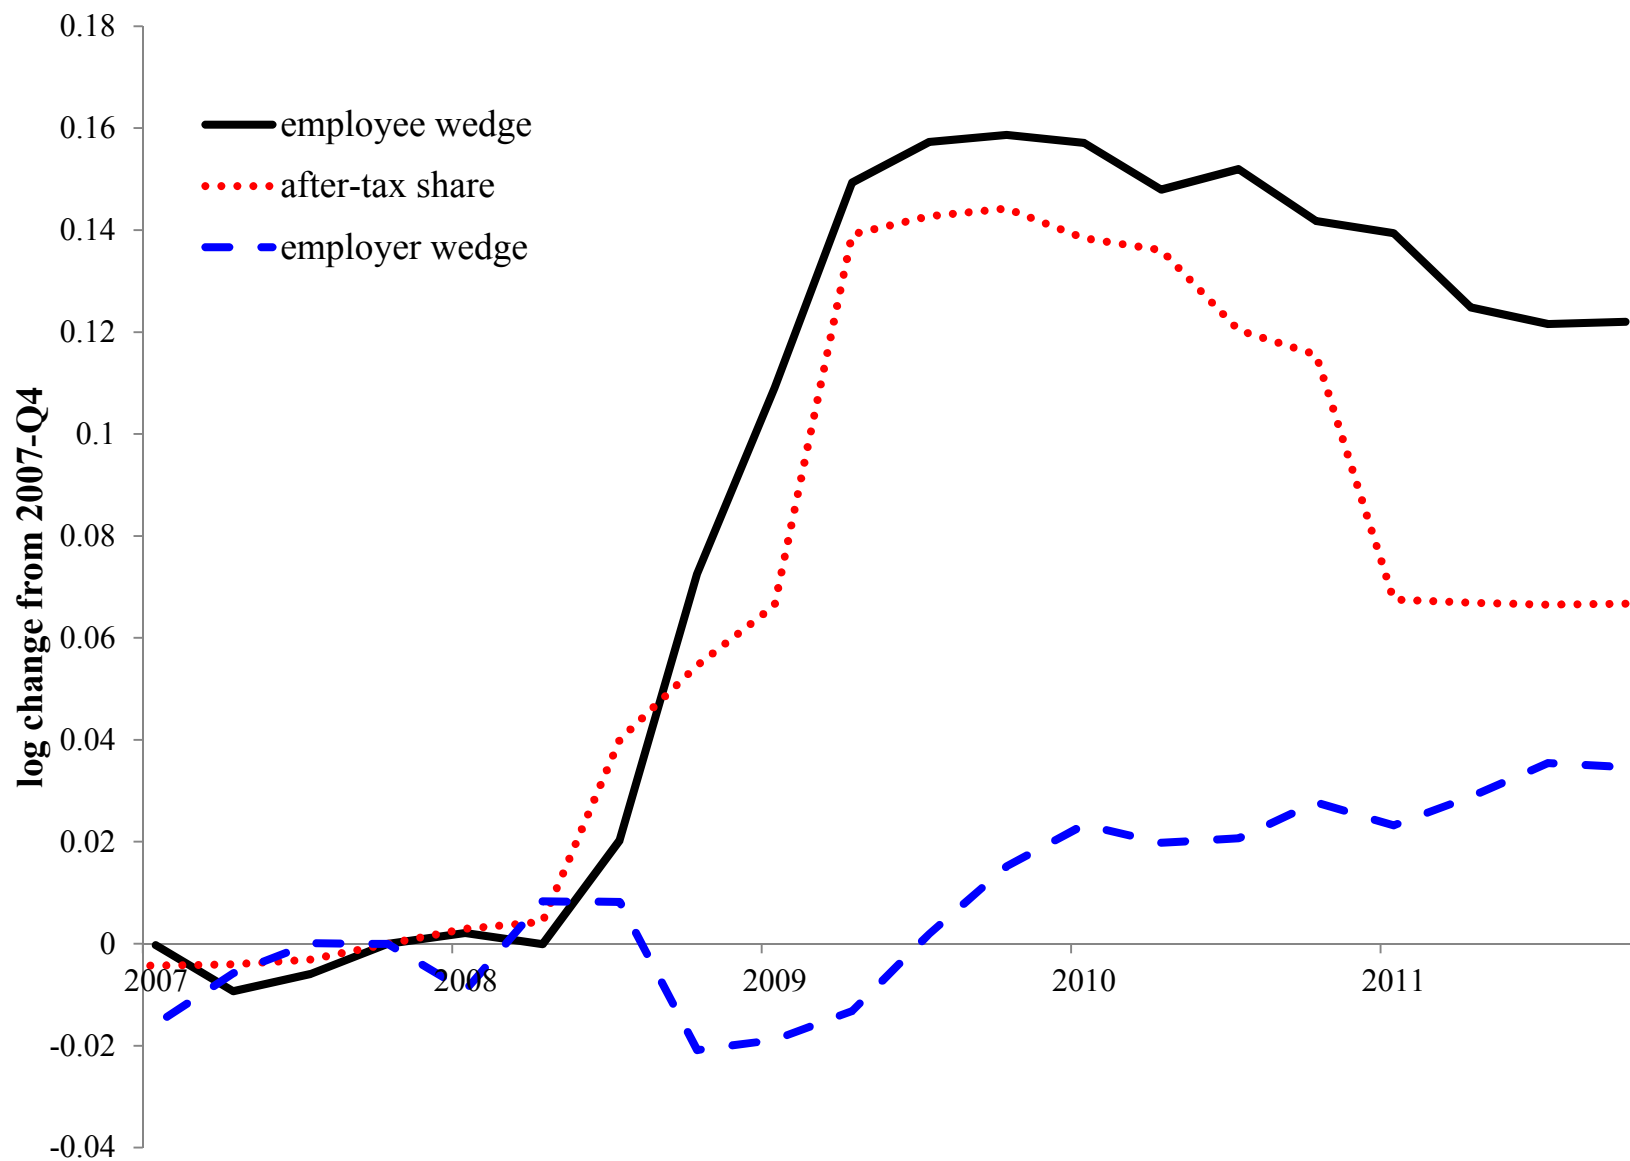


Figure 4. Statutory Marginal Labor Income Tax Rates
for non-elderly heads or spouses



Figure 5. Labor Market Wedges and Marginal Tax Rates



Bibliography

- Baker, Scott R., Nicholas Bloom, and Steven J. Davis. "Measuring Economic Policy Uncertainty." *manuscript, University of Chicago Booth School of Business.*, October 2011.
- Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen Terry. "Really Uncertain Business Cycles." *manuscript, Stanford University*, December 2011.
- Campbell, John Y., and John H. Cochrane. "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior." *Journal of Political Economy* 107, no. 2 (April 1999): 205-51.
- Carroll, Christopher D. "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Quarterly Journal of Economics* 112, no. 1 (February 1997): 1-55.
- Chicago Board Options Exchange. "The CBOE Volatility Index - VIX." 2009.
<http://www.cboe.com/micro/vix/vixwhite.pdf> (accessed January 6, 2012).
- Diamond, Douglas W., and Raghuram G. Rajan. "Fear of Fire Sales, Illiquidity Seeking, and Credit Freezes." *Quarterly Journal of Economics* 126, no. 2 (May 2011): 557-91.
- Eslami, Esa, Kai Filion, and Mark Strayer. "Characteristics of Supplemental Nutrition Assistance Program Households: Fiscal Year 2010." *Supplemental Nutrition Assistance Program (SNAP) Studies*. September 2011.
<http://www.fns.usda.gov/ora/menu/Published/snap/FILES/Participation/2010Characteristics.pdf> (accessed November 5, 2011).
- Garen, John E. "Executive Compensation and Principal-Agent Theory." *Journal of Political Economy* 102, no. 6 (1994): 1175-99.
- Holmstrom, Bengt, and Paul Milgrom. "Aggregation and Linearity in the Provision of Intertemporal Incentives." *Econometrica* 55, no. 2 (March 1987): 303-28.
- Lucas, Robert E. Jr. "Bernanke is the Best Stimulus Right Now." *Wall Street Journal*, December 2008.
- Meltzer, Allan H., and Scott F. Richard. "A Rational Theory of the Size of Government." *Journal of Political Economy* 89, no. 5 (October 1981): 914-27.
- Mirrlees, James A. "An Exploration in the Theory of Optimum Income Taxation." *Review of Economic Studies* 38, no. 114 (April 1971): 175-208.
- Mulligan, Casey B. "The ARRA: Some Unpleasant Welfare Arithmetic." *NBER working paper*, no. 18591 (February 2013).

- . *The Redistribution Recession*. New York: Oxford University Press (redistributionrecession.com), 2012.
- Neumann, John von, and Oskar Morgenstern. *Theory of Games and Economics Behavior*. Princeton, NJ: Princeton University Press, 1944.
- Peltzman, Sam. "The Growth of Government." *Journal of Law and Economics* 23, no. 2 (October 1980): 209-87.
- Perotti, Roberto. "Growth, Income Distribution, and Democracy: What the Data Say." *Journal of Economic Growth* 1, no. 2 (June 1996): 149-87.
- Romich, Jennifer L., Jennifer Simmelink, and Stephen Holt. "When Working Harder Does Not Pay: Low-Income Working Families, Tax Liabilities, and Benefit Reductions." *Families in Society: The Journal of Contemporary Social Services* 88, no. 3 (2007): 418-26.
- Rosen, Sherwin. "Authority, Control, and the Distribution of Earnings." *Bell Journal of Economics* 13 (Autumn 1982): 311-23.
- White House. *Franklin D. Roosevelt*. 2012.
<http://www.whitehouse.gov/about/presidents/franklindroosevelt> (accessed September 13, 2012).